

DETAILED ACTION

1. As per the instant Application having Application number 10/788,589, the examiner acknowledges the applicant's submission of the amendment 10/13/2009. At this point, no claims have been amended. Claims 1-2, 5-16, 26-28 and 31-36 are pending.
2. The Declaration filed on 10/13/2009 under 37 CFR 1.131 is sufficient to overcome the Hirakawa et al. (US 2005/0055523) reference.

REJECTIONS BASED ON PRIOR ART

Claim Rejections - 35 USC § 103

3. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

4. **Claims 1-2, 5-8, 11-16, 26-28 and 31-34** are rejected under 35 U.S.C. 103(a) as being unpatentable by Wu et al. (US 7,216,133) in view of Hirakawa et al. (US 2004/0267829).
5. As per claim 1, Wu discloses A method comprising: **[a method of accessing and controlling storage system depicted in (fig. 1) wherein replicas 102 and 104 are synchronized via logical view (col. 5, lines 64-66)]**
in response to receiving, from an application, a write request, **[updates or write transaction are made to replica 102 via inputs to tables 105 and 106 (col. 5, lines 42-51) wherein "synchronization is performed in response to an indicator received by replica 102 that**

synchronization should be initiated. The indicator may include: a recognition by replica 102 that it has items in tables 105 and 106 that have been changed or added through the user interface 134" (col. 6, lines 4-9)]

generating first and second write transactions, the first write transaction comprises a first tag, and the second write transaction comprises a second tag; [**"...the physical table 105 illustrated in FIG. 3A has two columns; namely, an item column 302 and a local change enumeration (LOCAL_CN) column 304. The item column 302 contains the items stored in the replica 102. The local change enumeration column 304 includes local change enumerations that the data store 126 assigns as labels to the items..." (col. 7, lines 37-58)** *(wherein items written to physical tables are interpreted as first write transaction having local change enumeration, interpreted as first tag)* in response to updates to physical tables in replica 102, logical tables 110 are updated to indicate updates and physical tables in replica 104 wherein "items are then sent from the logical view 110 in replica 102 to the logical view 112 in replica 104. Replica 104 can then map the sent items from the logical view 112 through the catalog 116 to the physical tables 107, 108, and 109" (col. 6, lines 4-17; figs. 1 and 3 and related text) *(thus generating a second write transaction comprising an item or data D)* wherein it is disclosed that "the physical tables 105 and 106; and 107, 108, and 109 may contain the same information" (col. 5, lines 31-42); thus, physical tables of replica 104 contain change enumerations *(interpreted as a second tag)*

transmitting the first and second write transactions to first and second storage devices, respectively; [updates in logical view 110 in replica 102 contains are mapped to physical tables 105 and 106 and updates in logical view 112 in replica 104 are mapped to physical

tables 107, 108 and 109 (fig. 1 and related text; col. 12, lines 52-62; col. 5, lines 25-65) which correspond to first and second storage devices wherein logical views are substantially similar to each other in different replicas to facilitate synchronization when updates are sent from a source replica to a destination replica (col. 11, lines 24-42)]

wherein the first write transaction further comprises data D to be written; wherein the second write transaction further comprises data D to be written; **[data items (col. 6, lines 18-40; figures 1 and 3A-3b and related text)]**

storing first write information in an entry of a first tag table, wherein the first write information comprises the first tag and an identity of a logical block of a first storage object where data D is to be written, wherein the first tag table is stored in first memory; **[updates written to replica 102 are stored in the combination of logical physical table and logical view synch table such as tables shown in figures 3A-3B and related text; wherein local_cn and version numbers correspond to a first tag wherein "the version in column 310 contains information about which replica made the change" (col. 8, lines 56-58) which corresponds to the identity of the logical view that made the change (See. Col. 6, lines 4-17). Wu further discloses synchronization/updating replicas at folder level wherein folders are logical constructs used for grouping items together (Col. 6, lines 18-40)]**

storing second write information in an entry of a second tag table, wherein the second write information comprises the second tag and an identity of a logical block of a second storage object where data D is to be written, wherein the second tag table is stored in second memory **[updates written to replica 104 are stored in the combination of logical physical table and logical view synch table as shown in figures 3A-3B and related text; wherein local_cn and**

version numbers correspond to a first tag wherein "the version in column 310 contains information about which replica made the change" (col. 8, lines 56-58) which corresponds to the identity of the logical view that made the change (Sec. Col. 6, lines 4-17). Wu further discloses synchronization/updating replicas at folder level wherein folders are logical constructs used for grouping items together (Col. 6, lines 18-40)]; however, Wu does not explicitly disclose each update comprises the identity of the logical block that where data is to be written.

Hirakawa discloses write transactions to first and second storage devices wherein write information such a write tags and identity of logical blocks wherein data is written is stored in first and second tables in first and second storage devices, respectively as **[each of storage systems 100a and 100b shown in figure 1 stores a journal as JNL1 and JNL2, respectively, wherein a "journal is composed of write data and update information. The update information is used to manage the write data and is composed of the time at which a write command was received, a group number an update number... the logical address of the write command (which corresponds to the claimed identify of logical blocks)... The update information may hold only one of either the time at which the write command was received or the update number (which corresponds to the claimed write tags)" (par. 0041) wherein "a retrieval method or a table for finding the logical address where the update information has been stored from the update number is provided in shared memory 140 of original storage system 100A" (par. 0154) "a retrieval method or a table for finding the logical address where the update information has been stored from the update number is provided in secondary storage system 100B" (par. 0156)].**

Wu and Hirakawa are analogous art in that they are of the same field of endeavor, that is, a system and/or method of memory control. Hirakawa suggests that it would have been desirable to incorporate information on write transactions such as tags and identity of logical blocks wherein this information is stored as journals in each storage device into the system of Wu because this would allow the system of Wu to **["transfer or replicate data between a plurality of storage systems, without affecting the host computer of the storage systems and also without affecting communication between the storage system and the host computer... to enable the reduction of data storage areas provided in a plurality of storage systems... to transfer or replicate data among a plurality of storage systems effectively, at high speed and in a seamless manner" (pars. 0008-0009)]**. Therefore, it would have been obvious to a person of ordinary skill in the art at the time the invention was made to modify Wu as suggested by Hirakawa to create a method to obtain the invention as specified in claim 1.

6. As per claim 2, the combination of Wu and Hirakawa discloses The method of claim 1 further comprising: generating third and fourth write transactions, wherein the third and fourth write transactions comprise third and fourth tags, respectively; transmitting the third and fourth write transactions to the first and second storage devices, respectively **[Wu discloses updates written to replicas 102 and 104 are stored in the combination of logical physical table and logical view synch table as shown in figures 3A-3B and related text; wherein local_cn and version numbers correspond to a first tag wherein "the version in column 310 contains information about which replica made the change" (col. 8, lines 56-58) which corresponds to the identity of the logical view that made the change (See. Col. 6, lines 4-17). Wu further**

teaches updates in logical view 110 in replica 102 contains are mapped to physical tables 105 and 106 and updates in logical view 112 in replica 104 are mapped to physical tables 107, 108 and 109 (fig. 1 and related text; col. 12, lines 52-62; col. 5, lines 25-65) which correspond to first and second storage devices wherein logical views are substantially similar to each other in different replicas to facilitate synchronization when updates are sent from a source replica to a destination replica (col. 11, lines 24-42)]. Hirakawa discloses [each of storage systems 100a and 100b shown in figure 1 stores a journal as JNL1 and JNL2, respectively, wherein a “journal is composed of write data and update information. The update information is used to manage the write data and is composed of the time at which a write command was received, a group number an update number... the logical address of the write command (*which corresponds to the claimed identify of logical blocks*)... The update information may hold only one of either the time at which the write command was received or the update number (*which corresponds to the claimed write tags*)” (par. 0041) wherein “a retrieval method or a table for finding the logical address where the update information has been stored from the update number is provided in shared memory 140 of original storage system 100A” (par. 0154) “a retrieval method or a table for finding the logical address where the update information has been stored from the update number is provided in secondary storage system 100B” (par. 0156)].

7. As per claim 5, the combination of Wu and Hirakawa discloses The method of claim 1 wherein: the first write transaction comprises data D to be written to a range of logical blocks of the first storage object; the second write transaction comprises data D to be written to a range of logical blocks of the second storage object [Wu discloses synchronization or updating replicas

at folder level wherein folders are logical constructs used for grouping items together (Col. 6, lines 18-40) wherein items are stored within folder boundaries and folders are maintained in replicas shown in figure 1 and updates to any replica are synchronized with the other. Hirakawa discloses each of storage systems 100a and 100b shown in figure 1 stores a journal as JNL1 and JNL2, respectively, wherein a “journal is composed of write data and update information. The update information is used to manage the write data and is composed of the time at which a write command was received, a group number an update number... the logical address of the write command (*which corresponds to the claimed identity of logical blocks*)... The update information may hold only one of either the time at which the write command was received or the update number (*which corresponds to the claimed write tags*)” (par. 0041) wherein “a retrieval method or a table for finding the logical address where the update information has been stored from the update number is provided in shared memory 140 of original storage system 100A” (par. 0154) “a retrieval method or a table for finding the logical address where the update information has been stored from the update number is provided in secondary storage system 100B” (par. 0156)].

8. As per claim 6, the combination of Wu and Hirakawa discloses The method of claim 5 wherein the first write information comprises an identity of the first storage object, and an identity of the range of logical blocks of the first storage object where data D is to be written, the second write information comprises, an identity of the second storage object, and an identity of the range of logical blocks in the second storage object where data D is to be written [Wu discloses items are identified by item id in col. 306 (Figs. 3A and 3B and related text) and

synchronization or updating replicas at folder level wherein folders are logical constructs used for grouping items together (Col. 6, lines 18-40) wherein items are stored within folder boundaries and folders are maintained in replicas shown in figure 1 and updates to any replica are synchronized with the other and also explains items in logical views are mapped to physical tables by catalogs in both replicas (Figs. 1, and 3A-3B and related text).

Hirakawa discloses each storage systems 200, 100a and 100b shown in figure 1 stores a journal wherein “each journal is comprised of write data and update information. The update information is information for managing write data, and comprises of the time at which a write command was received (update time), a group number, an update number in the group information 600 described later, a logical address of the write command (*which corresponds to the identify of logical blocks*), the size of the write data, and the logical address in the journal logical volume where the write data is stored. The update information may have only either the time at which the write command was received or the update number (*which corresponds to tags*)” (pars. 0075-0076; 0202 and 0204)].

9. As per claim 7, the combination of Wu and Hirakwadiscloses the method of claim 1 further comprising comparing the contents of one entry in the first tag table with the contents of entries in the second tag table to determine whether the second tag table includes an entry that matches the one entry [Wu discloses synchronization layer includes a “**synchronization change tracker that maintains versions and synchronization local change enumerations for the one or more items mapped to the logical view**” and wherein logical views are used to synchronize storage layers with one another; wherein local change enumerations and synchronization local change enumeration for different items are compared to determine

when to update a replica and wherein updates are sent from one replica to another via logical views (Figs. 3A and 3B and related text; col. 11, lines 16-52)].

10. As per claim 8, the combination of Wu and Hirakawa discloses The method of claim 7 further comprising copying data, associated with the logical block number identified by the one entry, from the first storage object to the logical block in the second storage object if the second table lacks an entry with contents matching the contents of the one entry [Wu discloses **“sending at least one mapped item from the source replica to a destination replica... replica 102 sending items from logical view 110 to logical view 112” (col. 11, lines 37-42). Hirakawa discloses using journal information for replicating data from one logical volume to another (pars. 0044-0048; fig. 1 and related text)].**

11. As per claim 11, the combination of Wu and Hirakawa discloses The method of claim 1 further comprising: generating a write transaction to write data to a logical block of a data volume; incrementing a counter in response to generating the write transaction; generating the first and second tags as a function of an output of the incremented counter [Wu discloses **sending data from one replica to another thru logical views wherein “the local change enumeration associated with each item is incrementally updated each time a change is made to an item. The data storage layer 126 assigns local change enumerations to all changes irrespective of whether those changes are made through the user interface or received through synchronization” (col. 7, lines 37-58; figs. 1, 3A-3B and related text; col. 11, lines 37-42). Hirakawa discloses “update information may hold only one of either the item at which a write command was received or the update number” (par. 0042)].**

12. As per claim 12, the combination of Wu and Hirakawa The method of claim 1 wherein the first and second storage devices comprise first and second object storage devices, respectively **[Wu discloses physical storage (col. 1, lines 16-51; col. 6, lines 18-40; fig. 1 and related text). Hirakawa discloses storage systems 100a and 100b (figure 1 and related text)]**.

13. As per claim 13, the combination of Wu and Hirakawa discloses The method of claim 1 wherein: the first write transaction comprises data D to be written to an extension of the first storage object; the second write transaction comprises data D to be written to an extension of the second storage object **[Wu discloses synchronization or updating replicas at folder level wherein folders are logical constructs used for grouping items together (Col. 6, lines 18-40) wherein items are stored within folder boundaries and folders are maintained in replicas shown in figure 1 and updates to any replica are synchronized with the other and also discloses change units and consistency units (col. 6, line 41-col. 7, line 36)]**.

14. As per claim 14, the combination of Wu and Hirakawa discloses The method of claim 13 further comprising: the first write information comprising an identity of the first storage object, and an indication that data D is to be stored in the extension of the first storage object, the second write information comprising an identity of the second storage object, and an indication that data D is to be stored in the extension of the second storage object **[Wu discloses synchronization or updating replicas at folder level wherein folders are logical constructs used for grouping items together (Col. 6, lines 18-40) wherein items are stored within folder boundaries and folders are maintained in replicas shown in figure 1 and updates to any replica are**

synchronized with the other and also discloses consistency units and change units (col. 6, line 41-col. 7, line 36)].

15. As per claim 15, the combination of Wu and Hirakawa discloses A method comprising: in response to receiving, from an application, a write request, generating first and second write transactions; wherein the first and second write transactions comprise first and second tags, respectively, wherein the first and second tags are identical to each other; transmitting the first and second write transactions to first and second storage devices, respectively; wherein the first write transaction comprises data D; wherein the second write transaction comprises data D; storing first write information in an entry of a first tag table, wherein the first write information comprises the first tag and an identity of a logical block of a first storage object where data D is to be written, wherein the first tag table is stored in first memory; storing second write information in an entry of a second tag table wherein the second write information comprises the second tag and an identity of a logical block of a second storage object where data D is to be written, wherein the second tag table is stored in second memory **[The rationale in the rejection to claim 1 is herein incorporated].**

16. As per claim 16, the combination of Wu and Hirakawa discloses The method of claim 1: wherein a computer system generates the first and second write transactions in response to generation of a write transaction by a first application executing on the computer system, wherein the first and second tags are generated by a first tag generator; a second computer system generating third and fourth transactions in response to generation of a write transaction by a second application executing on the second computer system; wherein the third and fourth write transactions comprise third and fourth tags, respectively, wherein the third and fourth tags

are generated by a second tag generator **[Wu discloses incrementally updating local change enumeration which may be received from a user interface or a result of synchronization (col. 7, lines 37-58) wherein user interface may use procedures or functions code (col. 5, lines 43-51)].**

17. As per claim 26, the combination of Wu and Hirakawa One or more computer readable mediums storing executable instructions executable, wherein a method is implemented in response to executing the instructions, the method comprising: in response to receiving a first transaction comprising a first tag and data D, storing first write information in an entry of a first tag table, wherein the first write information comprises the first tag and an identity of a logical block of a first storage object where the data D is to be written, wherein the first transaction is generated in response to a write request from an application; wherein the first tag table is stored in first memory; in response to receiving a second write transaction comprising a second tag and data D storing second write information in an entry of a second tag table, the second write information comprising the second tag and an identity of the logical block of a second storage object where data D is to be written, wherein the second tag table is stored in second memory, wherein the second transaction is generated in response to the write request from the application **[The rationale in the rejection to claim 1 is herein incorporated; further note that in Wu, a write transaction to replica 104 is generated in response to user input modifying data in replica 102 since synchronization is started in response to data being modified in replica 102 (See rejection to claim 1 above and Wu, col. 5, line 43-col. 6, line 17)].**

18. As per claim 27, the combination of Wu and Hirakawa discloses The method of claim 1 further comprising: synchronizing said first and second storage devices in response to an entry of

the first tag table not matching an entry of the second tag table [Wu discloses updates written to replicas 102 and 104 are stored in the combination of logical physical table and logical view synch table as shown in figures 3A-3B and related text; wherein local_cn and version numbers correspond to a first tag wherein "the version in column 310 contains information about which replica made the change" (col. 8, lines 56-58) which corresponds to the identity of the logical block/view that made the change (See. Col. 6, lines 4-17). Wu further teaches updates in logical view 110 in replica 102 contains are mapped to physical tables 105 and 106 and updates in logical view 112 in replica 104 are mapped to physical tables 107, 108 and 109 (fig. 1 and related text; col. 12, lines 52-62; col. 5, lines 25-65) wherein a "synchronization change tracker that maintains versions and synchronization local change enumerations for the one or more items mapped to the logical view" and wherein logical views are used to synchronize storage layers with one another; wherein local change enumerations and synchronization local change enumeration for different items are compared to determine when to update a replica and wherein updates are sent from one replica to another via logical views (Figs. 3A and 3B and related text; col. 11, lines 16-52)].

19. As per claim 28, the combination of Wu and Hirakawa discloses A method comprising: in response to receiving, from an application, a write request, generating first and second write transactions; generating first and second tags; transmitting the first write transaction to first storage device; transmitting the second write transaction to second storage device; transmitting the first tag to the first storage device; transmitting the second tag to the second storage device; wherein the first write transaction comprises data D to be written; wherein the second write transaction comprises data D to be written; receiving the first write transaction at the first storage

device; receiving the first tag at the first storage device; storing the first write information in an entry of a first tag table, wherein the write information comprises the first tag and an identity of a logical block of a first storage object where data D is to be written, wherein the first tag table is stored in first memory; receiving the second write transaction at the second storage device; receiving the second tag at the second storage device; storing the second write information in an entry of a second tag table, wherein the write information comprises the second tag and an identity of a logical block of a second storage object where data D is to be written, wherein the second tag table is stored in second memory [The rationale in the rejection to claim 1 is herein incorporated].

20. As per claim 31, the combination of Wu and Hirakawa discloses The one or more computer readable memories of claim 26 wherein: the first write transaction comprises data D to be written to a range of logical blocks of the first storage object; the second write transaction comprises data D to be written to a range of logical blocks of the second storage object [The rationale in the rejection to claim 5 is herein incorporated].

21. As per claim 32, the combination of Wu and Hirakawa discloses The one or more computer readable memories of claim 31 wherein: the first write information comprises an identity of the first storage object, and an identity of the range of logical blocks of the first storage object where data D is to be written; the second write information comprises an identity of the second storage object, and an identity of the range of logical blocks in the second storage object where data D is to be written [The rationale in the rejection to claim 6 is herein incorporated].

22. As per claim 33, the combination of Wu and Hirakawa discloses The one or more computer readable memories of claim 26 wherein the method further comprises comparing the contents of one entry in the first tag table with the contents of entries in the second tag table to determine whether the second tag table includes an entry that matches the one entry [**The rationale in the rejection to claim 7 is herein incorporated**].

23. As per claim 34, the combination of Wu and Hirakawa discloses The one or more computer readable memories of claim 33 wherein the method further comprises copying data, associated with the logical block number identified by the one entry, from the first storage object to the logical block in the second storage object if the second table lacks an entry with contents matching the contents of the one entry [**The rationale in the rejection to claim 8 is herein incorporated**].

24. **Claims 9-10 and 35-36** are rejected under 35 U.S.C. 103(a) as being unpatentable by Wu et al. (US 7,216,133) in view of Hirakawa et al. (US 2004/0267829) as applied to claims 1 and 7 above, and further in view of Crockett et al. (US 6,088,697).

25. As per claims 9-10 and 35-36, the combination of Wu and Hirakawa The method of claim 7, including [**Hirakawa discloses freeing updates having equal or less update numbers in journals, indicating the journal has been transmitted to the replica storage system (par. 0120, 0129)**] but does not disclose expressly deleting the one entry in the first table if the second table contains an entry with contents that match the contents of the one entry; deleting the entry in the second table with contents that match the contents of the one entry.

Crockett discloses [**a remote copy system in which a data structure maintains changes made to a primary data storage system and wherein the changes that have been**

applied to a secondary data storage volume are discarded from the data structure and the changes that have not been applied are maintained in the data structure (col. 1, lines 52-67)].

Wu, Hirakawa and Crockett are analogous art in that they are of the same field of endeavor, that is, a system and/or method of memory control. Crockett suggests that it would have been desirable to incorporate the ability to delete entries that have already been sent from one storage system to another from a table or data structure into the system of Wu and Hirakawa because this would provide improvement in performing remote copy operations for duplication of volumes from a primary site to a secondary site and minimizing clean-up time and time for synchronization or re-synchronizations [(col. 1, lines 52-67; col. 2, lines 60-64)]. Therefore, it would have been obvious to a person of ordinary skill in the art at the time the invention was made to modify Wu and Hirakawa as suggested by Crockett to incorporate the feature as claimed.

CLOSING COMMENTS

Examiner's Note

26. Examiner has cited particular columns and line numbers in the references as applied to the claims above for the convenience of the applicant. Although the specified citations are representative of the teachings in the art and are applied to the specific limitations within the individual claim, other passages and figures may apply as well. It is respectfully requested from the applicant, in preparing the responses, to fully consider the references in entirety as potentially

teaching all or part of the claimed invention, as well as the context of the passage as taught by the prior art or disclosed by the examiner.

When responding to this Office Action:

12. Applicant is requested to indicate where in the disclosure support is to be found for any new language added to the claims by amendment. 37 C.F.R. § 1.75(d)(1) requires such support in the Specification for any new language added to the claims and 37 C.F.R. § 1.83(a) requires support be found in the Drawings for all claimed features.

27. Applicant must clearly point out the patentable novelty which he or she thinks the claims present, in view of the state of the art disclosed by the references cited or the objections made, and must also explain how the amendments avoid the references or objections. See 37 C.F.R. § 1.111(c).

a. STATUS OF CLAIMS IN THE APPLICATION

28. The following is a summary of the treatment and status of all claims in the application as recommended by M.P.E.P. 707.07(i):

a(1) CLAIMS REJECTED IN THE APPLICATION

29. Per the instant office action, claims 1-2, 5-16, 26-28, and 31-36 have received an action on the merits and are subject of a non-final rejection.

a(2) CLAIMS NO LONGER UNDER CONSIDERATION

30. Claims 3-4, 17-25 and 29-30 have been canceled as of amendment dated 10/16/2009.

b. DIRECTION OF FUTURE CORRESPONDENCES

31. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Yaima Campos whose telephone number is (571) 272-1232. The examiner can normally be reached on Monday to Friday 8:30 AM to 5:00 PM.
32. If attempts to reach the above noted Examiner by telephone are unsuccessful, the Examiner's supervisor, Mr. Sanjiv Shah, can be reached at the following telephone number: Area Code (571) 272-4098.

The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300. Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

January 12, 2010

/Yaima Campos/
Examiner, Art Unit 2185